

Group Analysis of Self-organizing Maps based on Functional MRI using Restricted Frechet Means

Arnaud P. Fournel^{ab}, Emanuelle Reynaud^b, Michael J. Brammer^a,
Andrew Simmons^{ac}, and Cedric E. Ginestet^{ac}.

^aDepartment of Neuroimaging, Institute of Psychiatry, King's College London, UK, ^bLaboratoire d'Etude des Mécanismes Cognitifs (EMC), EA 3082, Université Lumière Lyon II, France,

^cNational Institute of Health Research (NIHR) Biomedical Research Centre for Mental Health.

Acknowledgments

This work was supported by a fellowship and core funds from the UK National Institute for Health Research (NIHR) Biomedical Research Centre for Mental Health (BRC-MH) at the South London and Maudsley NHS Foundation Trust and King's College London. This work has also been funded by the Guy's and St Thomas' Charitable Foundation as well as the South London and Maudsley Trustees. APF has received financial support from the Region Rhône-Alpes and the Université Lumière Lyon 2 through an ExploraDoc grant. We would also like to thank three anonymous reviewers for their valuable inputs.

Correspondence

Correspondence concerning this article should be sent to Cedric Ginestet at the Centre for Neuroimaging Sciences, NIHR Biomedical Research Centre, Institute of Psychiatry, Box P089, King's College London, De Crespigny Park, London, SE5 8AF, UK. Email may be sent to cedric.ginestet@kcl.ac.uk

Abstract

Studies of functional MRI data are increasingly concerned with the estimation of differences in spatio-temporal networks across groups of subjects or experimental conditions. Unsupervised clustering and independent component analysis (ICA) have been used to identify such spatio-temporal networks. While these approaches have been useful for estimating these networks at the subject-level, comparisons over groups or experimental conditions require further methodological development. In this paper, we tackle this problem by showing how self-organizing maps (SOMs) can be compared within a Frechean inferential framework. Here, we summarize the mean SOM in each group as a Frechet mean with respect to a metric on the space of SOMs. The advantage of this approach is twofold. Firstly, it allows the visualization of the mean SOM in each experimental condition. Secondly, this Frechean approach permits one to draw inference on group differences, using permutation of the group labels. We consider the use of different distance functions, and introduce two extensions of the classical sum of minimum distance (SMD) between two SOMs, which take into account the spatio-temporal pattern of the fMRI data. The validity of these methods is illustrated on synthetic data. Through these simulations, we show that the three distance functions of interest behave as expected, in the sense that the ones capturing temporal, spatial and spatio-temporal aspects of the SOMs are more likely to reach significance under simulated scenarios characterized by temporal, spatial and spatio-temporal differences, respectively. In addition, a re-analysis of a classical experiment on visually-triggered emotions demonstrates the usefulness of this methodology. In this study, the multivariate functional patterns typical of the subjects exposed to pleasant and unpleasant stimuli are found to be more similar than the ones of the subjects exposed to emotionally neutral stimuli. In this re-analysis, the group-level SOM output units with the smallest sample Jaccard indices were compared with standard GLM group-specific z -score maps, and provided considerable levels of agreement. Taken together, these results indicate that our proposed methods can cast new light on existing data by adopting a global analytical perspective on functional MRI paradigms.

KEYWORDS: Barycentre, Frechet Mean, fMRI, Group Comparison, Karcher mean, Multivariate analysis, Self-Organizing Maps, Unsupervised Learning.

Introduction

Self-organizing Maps (SOMs) were originally introduced by Teuvo Kohonen et al. (2000). A SOM is an unsupervised artificial neural network that describes a training data set as a (typically planar) layer of neurons or *output units*. Each neuron learns to become a prototype for a number of input units, until convergence of the algorithm. The resulting SOM therefore represents a projection of the inputs into a two-dimensional grid. In this sense, SOMs can be regarded as a dimension-reduction clustering algorithm. One of the main advantages of this unsupervised method is that the relative position of the neurons on the grid can be directly interpreted, in the sense that *proximity* of two units on the map indicates *similarity* of the prototypal profiles of these units.

SOMs have proved to be useful for data-driven analysis and have become popular tools in the machine learning community (Tarca et al., 2007). In neuroimaging, these methods have been successfully applied to the detection of fMRI response patterns related to different cognitive tasks (Liao et al., 2008, Ngan et al., 2002, Ngan and Hu, 1999, Wismüller et al., 2004). Since SOMs are non-parametric unsupervised neural networks, they do not require the specification of temporal signal profiles, such as haemodynamic response function or anatomical regions of interest in order to generate meaningful summaries of spatio-temporal patterns of brain activity. As a result, these methods have also been used to identify variations in low-frequency functional connectivity (Peltier et al., 2003). Statistically, however, observe that the absence of a probabilistic model can also be a limitation, as this does not allow for a formal evaluation of the goodness-of-fit of the method.

When used for clustering, the SOMs have the following main advantages. Firstly, starting with a sufficient number of neurons, the SOM procedure is able to identify features in the data even when these features are only typical of a small number of input vectors. Secondly, the resulting layer of neurons is arranged according to the similarity of these prototypes in the original data space. This ‘topology-preserving’ property is generally not available in other data-reduction techniques, such

as independent component analysis (ICA) or k -means clustering. This is one desirable property that SOMs share with multidimensional scaling. This topological structure facilitates the merging of nodes in order to form ‘superclusters’, which provide a way to visualize and compare high-dimensional fMRI data sets. Fischer and Hennig (1999), for instance, have demonstrated the specific relevance of these advantages to the analysis of experimental fMRI data.

One of the outstanding questions in the application of SOMs to fMRI data is whether one can summarize several subject-specific SOMs into a ‘mean map’, which would pool information over several subjects. In addition, it may be of interest to draw inference over group differences, by comparing the mean maps of several groups of subjects. Here, the term *group* is used interchangeably with the concept of experimental condition. Hence, two distinct groups need not be composed of different subjects, but may only represent different sets of measurements on the same individuals. One may, for instance, be interested in extracting the SOM that summarizes functional brain activity during a particular cognitive task; or in comparing the resting-state SOM signature of schizophrenic patients with that of normal subjects. Few studies, however, have tackled the problem of formally comparing two or more families of SOMs. Although several authors have proposed distance functions on spaces of SOMs (Kaski and Lagus, 1996, Deng, 2007, Kirt et al., 2007), to the best of our knowledge, none of these researchers have attempted to draw statistical inference on the basis of such comparisons. This lack of methodology highlights a pressing need for developing new strategies that would permit the extension of such multivariate methods from single-subject analysis to multiple group comparison.

SOM group analysis can naturally be articulated within a Frechean statistical framework. In 1948, Frechet introduced the concept of Frechet mean, sometimes referred to as barycentre or Karcher mean in the context of Euclidean and Riemannian geometry, respectively (Karcher, 1977). The Frechet mean extends this concept to any metric space. This quantity is a generalization of the traditional arithmetic mean, applied to abstract-valued random variables, defined over a metric space. The definition of a generalized notion of the arithmetic mean therefore solely relies on the specification of a metric on the data space of interest. Once such a metric has been specified, the Frechet mean is simply the element that minimizes a convex combination of the squared distances from all the elements in the space of interest. Hence, we can construct a metric space of SOMs by choosing a metric on that space, which permits the comparison of any two given SOMs in that space. Note that, in that context, the chosen pairwise distance function should be a proper metric in the sense that it should satisfy the four metric axioms: (i) non-negativity, (ii) coincidence, (iii) symmetry and (iv) the triangle inequality (see Searcoid, 2007, for an introduction to metric spaces). In the sequel, we consider the use of different distance functions on spaces of SOMs, which do not satisfy the triangle inequality. Nonetheless, we will show that such distance functions can easily be transformed into proper metrics, using a straightforward manipulation (Mannila and Eiter, 1997). See appendix A. The concept of the Frechet mean has proved to be useful in several domains of applications, including image analysis (Thorstensen et al., 2009, Bigot and Charlier, 2011), statistical shape analysis (Dryden and Mardia, 1998), and in the study of phylogenetic trees (Balding et al., 2009).

In this paper, our purpose is to use the concept of the Frechet mean for drawing statistical inference over several families of subject-specific SOMs. We thus construct Frechean independent and paired-sample t -statistics, by analogy with the classical treatment of real-valued random variables. Statistical inference for these different tests are then drawn using permutation of the group labels. In the paper at hand, these statistics will be constructed using the *restricted* Frechet mean, which has been shown to have desirable asymptotic properties (Sverdrup-Thygeson, 1981), but is more convenient to use from a computational perspective. The restricted Frechet mean is defined as the element in the sample space, which minimizes the squared distances from all the elements in the sample. This formal approach to group inference on families of subject-specific SOMs has the advantage of allowing a direct representation of the mean SOM in each group, thereby pooling together subject-specific information. In addition, the proposed methods also allow to formally draw inference at the group-level in terms of the chosen distance function.

The paper is organized as follows. In the next section, we give a general introduction to SOMs, and how they are computed highlighting the specific algorithm, which will be used throughout the

rest of the paper. In a third section, we describe our proposed Frechean framework for drawing inference on several groups of SOMs. This strategy is entirely reliant on the choice of metric for comparing two given SOMs, and we therefore dedicate a fourth section to the description of several distance functions on spaces of SOMs, which appear particularly well-suited for the analysis of fMRI data. These methods are tested on synthetic data, under a range of different conditions in section five, and on a real data set in section six. We close the paper by discussing the potential usefulness of this statistical strategy with an emphasis on the critical importance of the choice of the distance function.

Self-Organizing Maps (SOMs)

We assume here that an fMRI data set is available, which consists of several spatio-temporal volumes \mathbf{X}_i , with $i = 1, \dots, n_j$, for n_j subjects in the j^{th} experimental group. Each \mathbf{X}_i is a $V \times T$ matrix, with V voxels and T time points. In the sequel, it will be of interest to compare several families of such volumes, such that $j = 1, \dots, J$, for J experimental conditions. When describing the SOM inference algorithms, however, we will focus on a single subject-specific data set, \mathbf{X} .

Sequential algorithm

A SOM, denoted \mathbf{M} , consists of K output units or *neurons* arranged in a two-dimensional rectangular grid of size K where $K = k_1 \times k_2$. For convenience, we here assume that the grids of interest are square grids, such that $k_1 = k_2$. Thus, the units of a SOM will be indexed by $k = 1, \dots, K$, where k ‘reads’ the units in SOM from left to right and top to bottom. Each entry in \mathbf{M} is hence denoted by \mathbf{m}_k , and corresponds to the *coordinates* of that unit in \mathbf{M} . That is, \mathbf{m}_k is a two-dimensional vector representing the position of \mathbf{m}_k in \mathbf{M} , such that, for instance, $\mathbf{m}_1 = (1, 1)$, and $\mathbf{m}_2 = (1, 2)$, and so forth. Each output unit has an associated weight vector \mathbf{w}_k , which is, in our case, a time series over T data points.

The sequential SOM algorithm takes a set of V input units, \mathbf{x}_v ’s, corresponding to the rows of the input data \mathbf{X} . The steps of the procedure will be indexed by $\gamma = 0, \dots, \Gamma$, which denote the iterations of the algorithm, and Γ is the final step at which a stopping condition is satisfied. In our case, the stopping rule is simply the number of iterations, but more sophisticated convergence-based criteria can be used. We firstly initialize the output units in \mathbf{M} as random draws from a uniform distribution on \mathbb{R}^T . Secondly, an input vector, denoted \mathbf{x}_v , is randomly chosen amongst the V time series. All V voxels are selected at each step of the algorithm, and these input vectors are therefore dependent on γ . We will thus denote this dependence on the iterations by $\mathbf{x}_v(\gamma)$. For each input vector presented to \mathbf{M} , we identify the unit in \mathbf{M} , which is the ‘closest’ to the input $\mathbf{x}_v(\gamma)$. Here, closeness is generally measured in terms of Euclidean distance with respect to the values taken by the input vectors. The unit in \mathbf{M} , which is the closest to $\mathbf{x}_v(\gamma)$ is referred to as the Best Matching Unit (BMU). The *index* of that BMU, for a given input vector \mathbf{x}_v at iteration γ , is defined as follows,

$$c(v, \gamma) = \underset{k \in \{1, \dots, K\}}{\operatorname{argmin}} \|\mathbf{x}_v(\gamma) - \mathbf{w}_k\|, \quad (1)$$

with $\|\cdot\|$ denoting the Euclidean norm on \mathbb{R}^T . Here, $\mathbf{x}_v(\gamma)$ and \mathbf{w}_k are T -dimensional time series. Thirdly, we update the BMU and its neighbors. The new values of these units are defined as a linear relationship of the input vector $\mathbf{x}_v(\gamma)$. For a given $\mathbf{x}_v(\gamma)$, the updating rule for the BMU and its neighbors is the following,

$$\mathbf{w}_k(\gamma + 1) = \mathbf{w}_k(\gamma) + \alpha(\gamma) K_\gamma(\mathbf{m}_k, \mathbf{m}_{c(v, \gamma)}) (\mathbf{x}_v(\gamma) - \mathbf{w}_k(\gamma)), \quad (2)$$

for every $k = 1, \dots, K$. After updating their weights, the BMU and its neighbours are closer to $\mathbf{x}_v(\gamma)$ in the sense that they constitute a better representation of that input vector. These steps are repeated for a fixed number of iterations, Γ . The updating rule in equation (2) contains two key parameters: (i) the learning rate, denoted $\alpha(\gamma)$ and (ii) the kernel function represented by

$K_\gamma(\mathbf{m}_k, \mathbf{m}_{c(v,\gamma)})$, which grows smaller as we consider units in \mathbf{M} , which are further away from the BMU in the space of coordinates of \mathbf{M} . We describe these two quantities in turn.

The learning rate, $\alpha(\gamma)$ in equation (2), is a decreasing function of the number of iterations, γ , which controls the amount of learning accomplished by the algorithm –that is, the dependence of the values of the units in \mathbf{M} on the inputs. By convention, we have $\alpha(\gamma) \in [0, 1]$ for every γ . Three common choices for $\alpha(\cdot)$ are a linear function, a function inversely proportional to the number of iterations and a power function, such as the following recursive definition,

$$\alpha(\gamma + 1) = \left(\frac{\alpha(0)}{\alpha(\gamma)} \right)^{\gamma/\Gamma}, \quad (3)$$

for every $\gamma = 1, \dots, \Gamma$. A popular initialization for the learning rate is $\alpha(0) = 0.1$ (Peltier et al., 2003, González and Dasgupta, 2003). Clearly, as the algorithm progresses towards Γ , the value of $\alpha(\gamma)$ decreases towards 0. Note, however, that, in the paper at hand, we use the batch version of this algorithm, which does not require the specification of a learning rate.

In equation (2), we have also made use of the neighborhood kernel, $K_\gamma(\mathbf{m}_k, \mathbf{m}_{c(v,\gamma)})$. As for the learning rate, the value taken by this kernel decreases with the number of iterations, and is therefore dependent on γ . This dependence on γ has been emphasized through a subscript on K . For a given output unit $c(v, \gamma)$ in the map \mathbf{M} , the neighborhood kernel, $K_\gamma(\mathbf{m}_k, \mathbf{m}_{c(v,\gamma)})$, quantifies how ‘close’ is \mathbf{m}_k to the BMU, which has index $c(v, \gamma)$. Observe that this closeness is expressed in terms of Euclidean distances on the grid *coordinates*. As commonly done in this field, we here choose a standard Gaussian kernel to formalize the dependence of each unit on the values of its neighbors, such that

$$K_\gamma(\mathbf{m}_k, \mathbf{m}_{c(v,\gamma)}) = \exp \left(-\frac{\|\mathbf{m}_k - \mathbf{m}_{c(v,\gamma)}\|^2}{2\sigma(\gamma)^2} \right), \quad (4)$$

where $\|\cdot\|^2$ represents the two-dimensional dot product. Here, $\sigma(\gamma)$ is a linear function of the number of iterations, which controls the size of the neighborhood around the BMU. This function is defined recursively as $\sigma(\gamma + 1) = \sigma(0)(1 - \gamma/\Gamma)$, where $\sigma(0)$ is a parameter value that represents the initial neighborhood radius. This parameter is commonly initialized with respect to the size of the two-dimensional grid, \mathbf{M} , such that $\sigma(0) = k_1$, which is the ‘height’ of the output SOM.

Batch algorithm

A popular alternative to the sequential SOM algorithm described in the previous section is the *batch SOM algorithm*, which has the advantage of being more computationally efficient than its sequential counterpart (Vesanto and Alhoniemi, 2000). It has been successfully used in the context of fMRI analysis (Ngan et al., 2002), in natural language processing (Kohonen et al., 2000), and in the face recognition literature (Tan et al., 2005). The main difference between these two approaches is that the entire training set is considered at once in the batch SOM algorithm, which permits the updating of the target SOM with the net effect of all the inputs.

This ‘global’ updating is performed by replacing the input vector, denoted $\mathbf{x}_v(\gamma)$ in the previous section, with a weighted average of the input vectors, where the relative weight of each input vector is proportional to the neighborhood kernel values. At the γ^{th} step of the algorithm, we are therefore conducting the following global updating,

$$\mathbf{w}_k(\gamma + 1) = \frac{\sum_{v=1}^V \mathbf{x}_v K_\gamma(\mathbf{m}_k, \mathbf{m}_{c(v,\gamma)})}{\sum_{v=1}^V K_\gamma(\mathbf{m}_k, \mathbf{m}_{c(v,\gamma)})}, \quad (5)$$

for every $k = 1, \dots, K$. It can easily be seen from equation (5) that $\mathbf{w}_k(\gamma + 1)$ is a convex linear combination of the input vectors, \mathbf{x}_v ’s, where each of the V inputs is weighted by $K_\gamma(\mathbf{m}_k, \mathbf{m}_{c(v,\gamma)}) / \sum_{v=1}^V K_\gamma(\mathbf{m}_k, \mathbf{m}_{c(v,\gamma)})$, and the sum of these weights is equal to 1. Another non-negligible advantage of the batch SOM algorithm is that it removes the dependence of the outputs on the learning rate parameter, denoted $\alpha(\gamma)$, as stated in the previous section.

Throughout the rest of the paper, we will make use of the batch algorithm, with $\sigma(0) = k_1$, and $k_1 = k_2 = 3$, thereby producing SOMs of dimensions 3×3 . Output units in all SOMs are initialized randomly. Other groups of researchers in neuroimaging have used square SOMs (Peltier et al., 2003, Liao et al., 2008). However, we have also investigated using simulated data, whether the specification of rectangular maps had a significant impact on our proposed inferential methods (see appendix B).

SOM Group Frechean Inference

The question of inferring the statistical significance of the difference between two families of SOMs can be addressed through the use of abstract-valued random variables as advocated by Fréchet (1948). In this approach, random variables are solely defined with respect to a probability measure on a metric space, (\mathcal{X}, d) (see Parthasarathy, 1967, chap. 2). Hence, it suffices to define a metric on the space of interest, in order to obtain a valid statistical framework. Once such a metric has been chosen, one can construct the mean element in that space, which is commonly referred to as the Frechet mean.

In the paper at hand, we are considering a space of SOMs, which we may denote by (\mathcal{M}, d) , where d is a metric on that space. A range of different distance functions for such spaces of SOMs will be described in the next section. As in most standard fMRI designs, we assume that we have J experimental conditions, with n_j subjects in each condition, thereby allowing for a different number of subjects in each experimental condition. A full data set will be summarized as an array of SOMs, $\{\mathbf{M}_{ij}\}$, with $i = 1, \dots, n_j$ and $j = 1, \dots, J$, such that \mathbf{M}_{ij} corresponds to the SOM of the i^{th} subject in the j^{th} condition. Given such a sample of SOMs, we can then define the Frechet mean for the j^{th} condition as follows,

$$\widehat{\mathbf{M}}_j = \operatorname{argmin}_{\mathbf{M}' \in \mathcal{M}} \frac{1}{(n_j - 1)} \sum_{i=1}^{n_j} d(\mathbf{M}_{ij}, \mathbf{M}')^2, \quad (6)$$

where we have used the Bessel's correction (i.e. $n_j - 1$) by analogy with the real-valued setting. Given the complexity of the underlying space of SOMs, such a minimization may be unwieldy. As a result, it is computationally more practical to consider the *restricted* Frechet mean, as introduced by Sverdrup-Thygeson (1981). The classical Frechet mean in equation (6) is obtained by identifying the element in the population of SOMs, which minimizes the average squared distances from all the elements in the sample. The restricted Frechet mean, by contrast, is obtained by identifying the element in the *sample*, which has this property. Hence, the restricted Frechet mean is computed as follows,

$$\overline{\mathbf{M}}_j = \operatorname{argmin}_{\mathbf{M}' \in \mathbf{\Lambda}_j} \frac{1}{(n_j - 1)} \sum_{i=1}^{n_j} d(\mathbf{M}_{ij}, \mathbf{M}')^2, \quad (7)$$

where $\mathbf{\Lambda}_j$ denotes the sampled n_j SOMs in the j^{th} condition, such that $\mathbf{\Lambda}_j = \{\mathbf{M}_{1,j}, \dots, \mathbf{M}_{n_j,j}\}$. The restricted Frechet mean has been shown to be consistent, through a generalization of the strong law of large numbers due to Sverdrup-Thygeson (1981). Asymptotically, $\overline{\mathbf{M}}_j$ converges almost surely to a subset of the theoretical restricted mean, which takes values in the support of the target population distribution. In the sequel, the theoretical restricted Frechet mean for the j^{th} condition will be denoted by μ_j , following standard convention.

Similarly, one can define the condition-specific sample Frechet variances. These quantities are simply the values taken by the criteria, which are minimized in equation (7), such that the (restricted) Frechet variance for the j^{th} condition is defined with respect to the restricted Frechet mean in the following manner, for every $j = 1, \dots, J$,

$$S_j^2 = \frac{1}{(n_j - 1)} \sum_{i=1}^{n_j} d(\mathbf{M}_{ij}, \overline{\mathbf{M}}_j)^2. \quad (8)$$

Using the restricted Frechet mean and variance, it is now possible to construct a non-parametric t -test on the metric space of SOMs. Here, we therefore assume that we solely have two experimental conditions, such that $J = 2$. The null hypothesis stating that the (restricted) Frechet means of these two distributions are δ_0 -separated, can be formally expressed as follows, $H_0 : d(\mu_1, \mu_2) = \delta_0$. Naturally, our interest will especially lie in testing the null hypothesis stating that there is no difference between the theoretical restricted Frechet means, which corresponds to $H_0 : d(\mu_1, \mu_2) = 0$. This can be tested using the following Frechet t -statistic,

$$t_F = \frac{d(\overline{\mathbf{M}}_1, \overline{\mathbf{M}}_2) - \delta_0}{S_p (1/n_1 + 1/n_2)^{1/2}}, \quad (9)$$

where the denominator, S_p , is the classical pooled sample variance, which is defined by analogy with the real-valued setting as

$$S_p^2 = \frac{(n_1 - 1)S_1^2 + (n_2 - 1)S_2^2}{n_1 + n_2 - 2}.$$

In addition, if one is considering two samples of equal sizes and assuming equal Frechet variances, then the aforementioned t_F -statistic for such a mean difference can be defined as follows

$$t_F = \frac{d(\overline{\mathbf{M}}_1, \overline{\mathbf{M}}_2)}{S_p / \sqrt{N}}, \quad (10)$$

where, in this case, the pooled variance is simply the sum of the variances of the two samples, such that $S_p^2 = S_1^2 + S_2^2$, and with $N = n_1 + n_2$. Statistical inference is then conducted using permutation on the group labels. Although our proposed t_F -statistic is a real-valued random variable, its asymptotic distribution is unknown. Indeed, the behavior of this statistic depends on a large number of other random variables, which are combined using the non-linear procedure for obtaining group-level SOMs. As a result, the permutation-based distribution of t_F under the null hypothesis is not expected to follow a standard t -distribution. In particular, the null distribution of t_F need not be symmetric. Since we are here solely considering a generalization of the t -test, but will be applying this statistic to more than two experimental conditions, we will also make use of the standard Bonferroni correction for multiple testing.

Choice of Distance Functions

In our proposed approach to group comparison, the choice of the metric on the space of SOMs is paramount. Different distance functions capture different aspects of the SOMs under scrutiny. It is therefore of interest to evaluate group differences with respect to several choices of distance functions. We here review the main distance functions, which have been previously proposed in the literature for comparing two given SOMs. In addition, we introduce a spatio-temporal sum of minimum distances, which is especially relevant for the study of fMRI-based SOMs.

Quantization Error and Other Measures

This measure is not a metric, but a popular tool for evaluating the accuracy of the SOM generated from a given data set. The so-called quantization error measures the average quantization error of the target SOM (Kohonen, 2001). It is defined as the sum of the Euclidean distances between each input unit, \mathbf{x}_v , and its best matching prototype on \mathbf{M} —that is, the BMU of \mathbf{x}_v . The quantization error, denoted Q_e , is thus formally defined as follows,

$$Q_e(\mathbf{M}, \mathbf{X}) = \sum_{v=1}^V \|\mathbf{x}_v - \mathbf{w}_{c(v)}\|,$$

where, as before, $\|\cdot\|$ denotes the T -dimensional Euclidean norm and $c(v)$ is the index of the BMU in \mathbf{M} with respect to \mathbf{x}_v as described in equation (1). This measure is a good indicator of the

convergence of a SOM, and is often used when assessing the behavior of the algorithms described in the previous section. In this paper, we will use a variant of the quantization error in order to identify the output units, which explain the largest amount of between-subject ‘variance’ in the data. However, the quantization error does not allow the computation of the distance between two given SOMs.

Kaski and Lagus (1996) have proposed a measure of dissimilarity between two SOMs. They proceeded by comparing the shortest path on each SOM after matching a given pair of input vectors. This dissimilarity measure is computed by comparing the distances between all pairs of data samples on the feature maps. This method, however, is not computationally efficient, and would be especially challenging when considering fMRI data sets, where neuroscientists are commonly handling about 100,000 input vectors –that is, the voxel-specific time series– for every subject.

Sum of Minimum Distances (T-SMD)

The Sum of Minimum Distances (SMD) was originally introduced by Mannila and Eiter (1997) and has been widely used in image recognition and retrieval (Kriegel, 2004, Takala et al., 2005, Tungaraza et al., 2009). Moreover, the SMD function and some of its variants have already been used in order to tackle the problem of comparing several SOMs (Deng, 2007). Given two SOMs, denoted \mathbf{M}_x and \mathbf{M}_y for input data sets \mathbf{X} and \mathbf{Y} , respectively, the SMD can be computed as follows. For every unit, \mathbf{w}_x in \mathbf{M}_x , we calculate the Euclidean distance between \mathbf{w}_x and every unit \mathbf{w}_y in \mathbf{M}_y in order to retain the unit in \mathbf{M}_y that minimizes this distance. These minimal distances are summed and then normalized by the total number of input vectors, denoted V , in our case. This gives an \mathbf{M}_x -to- \mathbf{M}_y score. The same procedure is performed in the opposite direction in order to produce an \mathbf{M}_y -to- \mathbf{M}_x score. The average of the \mathbf{M}_x -to- \mathbf{M}_y and the \mathbf{M}_y -to- \mathbf{M}_x scores is then defined as the overall SMD between \mathbf{M}_x and \mathbf{M}_y . Therefore, this distance function compares SOMs on the basis of the dissimilarity of the time series underlying each output unit. It follows that this procedure mainly emphasizes temporal differences between the fMRI volumes of interest. Thus, we will label this classical SMD as *temporal SMD*, and denote it by T-SMD. It is formally defined as

$$\text{T-SMD}(\mathbf{M}_x, \mathbf{M}_y) = \frac{1}{2V} \left(\sum_{\mathbf{w}_x \in \mathbf{M}_x} \min_{\mathbf{w}_y \in \mathbf{M}_y} d_e(\mathbf{w}_x, \mathbf{w}_y) + \sum_{\mathbf{w}_y \in \mathbf{M}_y} \min_{\mathbf{w}_x \in \mathbf{M}_x} d_e(\mathbf{w}_y, \mathbf{w}_x) \right), \quad (11)$$

where $d_e(\mathbf{w}_x, \mathbf{w}_y) = \|\mathbf{w}_x - \mathbf{w}_y\|$ is the Euclidean distance between \mathbf{w}_x and \mathbf{w}_y on \mathbb{R}^T , where \mathbf{w}_x and \mathbf{w}_y represent T -dimensional prototypical time series for maps \mathbf{M}_x and \mathbf{M}_y , respectively.

It is important to note that the SMD function can be re-written by treating a map, \mathbf{M}_x , as a set of weight vectors, \mathbf{w}_x . In this case, we consider the metric space of all weight vectors, \mathbf{w}_x . This metric space is (\mathbb{R}^T, d_e) . By a slight abuse of notation, the SOM, \mathbf{M}_x , will be used to denote the set of all output vectors, \mathbf{w}_x associated with the units in \mathbf{M}_x . Therefore, we have $\mathbf{M}_x \subset \mathbb{R}^T$. As a result, we can apply the classical definition of the distance between the subset of a metric space and an element of that space, $\mathbf{w}_x \in \mathbb{R}^T$, such that $d(\mathbf{w}_x, \mathbf{M}_x) = \min\{d(\mathbf{w}_x, \mathbf{w}'_x) : \mathbf{w}'_x \in \mathbf{M}_x\}$. Using these conventions, it becomes possible to reformulate the SMD function in equation (11) in the following manner as stated by Mannila and Eiter (1997),

$$\text{T-SMD}(\mathbf{M}_x, \mathbf{M}_y) = \frac{1}{2V} \left(\sum_{\mathbf{w}_x \in \mathbf{M}_x} d_e(\mathbf{w}_x, \mathbf{M}_y) + \sum_{\mathbf{w}_y \in \mathbf{M}_y} d_e(\mathbf{w}_y, \mathbf{M}_x) \right).$$

In addition, observe that the SMD function is not in general a proper metric, in the sense that the triangle inequality may fail to be satisfied (see Mannila and Eiter, 1997, for a counterexample). However, one can easily produce a proper metric through the identification of the shortest paths between any two elements in the space of interest, and then define a new metric with respect to these shortest paths (see appendix A). It can easily be shown that such a transformation necessarily produces proper metrics, when considering metrics based on the SMD function (Mannila and Eiter,

1997). This particular procedure can easily be implemented in our case, because we have focused our attention on the restricted Frechet mean, where the minimization required to identify the mean element is solely conducted over the space of the sampled elements. As a result, there exists a small number of possible shortest paths between every pair of elements in the sample, which greatly facilitates the required transformation for producing a proper metric. This procedure was systematically conducted in the sequel, and therefore all the variants of the SMD function utilized in this paper are indeed proper metrics. We will thus assume throughout this paper that all distance functions have been adequately transformed. We now introduce two novel variants of the SMD function, which take into account the spatial and spatio-temporal properties of the fMRI data.

Spatial SMD

One may also be interested in quantifying the amount of ‘spatial overlap’ between two given SOMs. This question is especially pertinent when analyzing SOMs based on fMRI data sets. Here, we therefore wish to evaluate whether the units in two different maps correspond to similar subsets of voxels in the original images. Such a distance can be quantified through a slight modification of the aforementioned SMD metric, where the Hamming distance is used in the place of the Euclidean distance.

$$\begin{aligned} \text{S-SMD}(\mathbf{M}_x, \mathbf{M}_y) = \frac{1}{2V} & \left(\sum_{\mathbf{w}_x \in \mathbf{M}_x} \min_{\mathbf{w}_y \in \mathbf{M}_y} \text{Ham}(S(\mathbf{w}_x), S(\mathbf{w}_y)) \right. \\ & \left. + \sum_{\mathbf{w}_y \in \mathbf{M}_y} \min_{\mathbf{w}_x \in \mathbf{M}_x} \text{Ham}(S(\mathbf{w}_y), S(\mathbf{w}_x)) \right). \end{aligned} \quad (12)$$

with V denoting the number of voxels in the fMRI volumes of interest, and where $S(\mathbf{w}_x)$ denotes the binarized index vector of the voxels, which have been assigned to unit \mathbf{w}_x in \mathbf{M}_x . That is, if the voxel v has been assigned to \mathbf{w}_x , then $S_v(\mathbf{w}_x) = 1$, otherwise, we have $S_v(\mathbf{w}_x) = 0$. In addition, we have here made use of the celebrated Hamming distance, which takes the following form (Hamming, 1950),

$$\text{Ham}(S(\mathbf{w}_x), S(\mathbf{w}_y)) = \frac{1}{V} \sum_{v=1}^V \mathcal{I}\{S_v(\mathbf{w}_x) \neq S_v(\mathbf{w}_y)\}, \quad (13)$$

where $\mathcal{I}\{f(x)\}$ stands for the indicator function taking a value of 1 if $f(x)$ is true, and 0 otherwise. Here, the term spatial refers to the spatial distribution of the voxels allocated to a particular output unit. Hence, S-SMD does not emphasize the spatial location of the output units, as these allocations are arbitrary, but rather the spatial distribution of the voxels allocated to that output unit. Observe that we have here solely considered differences in the spatial distributions of the best matched pair of SOM units, where that matching is done through the minimization reported in equation (12). This approach, however, omits to take into account the similarity of these SOM units as prototypal time series. Both the spatial and the temporal aspects of these maps can nonetheless be combined, as described in our proposed spatio-temporal SMD.

Spatio-Temporal SMD

In this novel variant of the classical SMD, we are quantifying the amount of spatial overlap between any pair of output units in two distinct maps. In contrast to the S-SMD described in the previous section, however, we are here comparing the spatial distributions (i.e. the sets of voxel indexes assigned to a particular unit) of the units that are the closest in terms of time series profiles, thereby combining the temporal and spatial properties of the data. This spatio-temporal version

of the SMD function is defined as follows,

$$\text{ST-SMD}(\mathbf{M}_x, \mathbf{M}_y) = \frac{1}{2} \left(\sum_{\mathbf{w}_x \in \mathbf{M}_x} \text{Ham}(S(\mathbf{w}_x), S(\mathbf{w}_y^*)) + \sum_{\mathbf{w}_y \in \mathbf{M}_y} \text{Ham}(S(\mathbf{w}_y), S(\mathbf{w}_x^*)) \right), \quad (14)$$

where

$$\mathbf{w}_y^* = \underset{\mathbf{w}_y \in \mathbf{M}_y}{\text{argmin}} d_e(\mathbf{w}_x, \mathbf{w}_y), \quad \text{and} \quad \mathbf{w}_x^* = \underset{\mathbf{w}_x \in \mathbf{M}_x}{\text{argmin}} d_e(\mathbf{w}_y, \mathbf{w}_x);$$

where, again, $d_e(\mathbf{w}_x, \mathbf{w}_y) = \|\mathbf{w}_x - \mathbf{w}_y\|$ is the Euclidean distance on \mathbb{R}^T , and $S(\mathbf{w}_x)$ is the index set of the voxels in \mathbf{X} , whose time series are best represented by \mathbf{w}_x . Albeit the formulae in equation (14) is somewhat convoluted, it corresponds to, perhaps, the most intuitive perspective on the problem of comparing SOMs, when using fMRI data. Indeed, we are quantifying the amount of spatial overlap between units, which are as similar as possible in terms of their temporal profiles.

We summarize this section with a concise description of these three different types of SOM distance functions:

- i. *Temporal SMD* (T-SMD) is based on the sum of the minimum Euclidean distances between the time series of the output units.
- ii. *Spatial SMD* (S-SMD) is based the sum of the minimum Hamming distances between the sets of voxels allocated to the output units.
- iii. *Spatio-temporal SMD* (ST-SMD) is based on the sum of Hamming distances between the sets of voxels allocated to the output units, which are the most similar in terms of their time series.

Synthetic Data Simulations

The proposed methods were tested on three different simulated scenarios, with varying degrees of difficulty. In particular, by isolating different types of differences between the two groups of interest, each of these scenarios emphasizes the need for the use of specific metrics, capturing different aspects of the spatio-temporal process under investigation. In this sense, our simulations strive to produce realistic differences between the two families of fMRI volumes, where group differences may be either spatial, temporal or both spatial and temporal.

Simulation Scenarios

For each simulated data set we have constructed two groups of 20 subjects, where each subject-specific data set is composed of two-dimensional images with 10×10 voxels, over 50 time points, as represented in panels (a-c) of figure 1. The time series at each voxel can be of three different types, as illustrated in panels (d-f) of figure 1, composed of two different signals and one background time series. The two signals represented in panels (d) and (e) are sinusoids oscillating over $[-1, 1]$, with a frequency of 1/10Hz and 1/20Hz, respectively. We have then added a vector of Gaussian random noise, \mathbf{z} , to these two types of time series, such that $z_t \sim N(0, \sigma^2)$, for every $t = 1, \dots, 50$, and for different choices of σ . The background noise time series, in panel (f), is solely composed of the random noise for a given σ .

The three scenarios in panels (a-c) of figure 1 are ordered in terms of the degree of ‘separability’ of the two groups, where the easiest scenario is on the left and the most difficult one is on the right. The first data set (SC1) was built with three different time series at three different locations, corresponding to the purple, blue and gray colors. In this scenario, the groups differ both in terms of the temporal profiles of some of their voxels and in terms of the spatial locations of these voxels. The second scenario (SC2) was constructed with two different types of time series. Here, the two groups solely differ in terms of the temporal profiles of some of their voxels. Finally, in the third

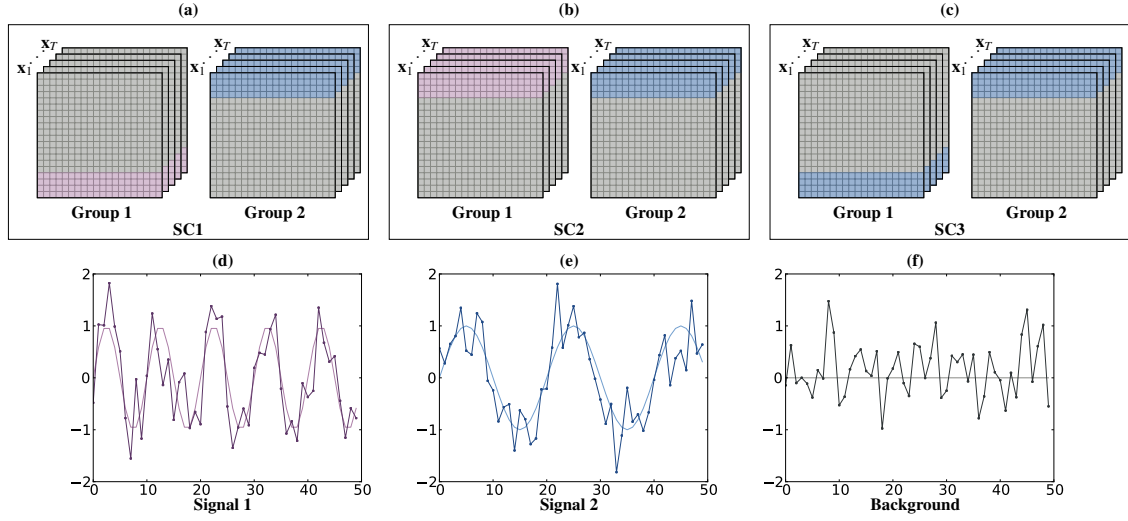


Figure 1. Description of the three simulated scenarios ordered by increasing levels of difficulty. In panels (a-c), we have reported the spatial distributions of the input vectors for each scenario, where each data set is composed of (10×10) -images over $T = 50$ time points. In panels (d-f), we have represented the three types of time series used in these simulations. Here, SC1, SC2 and SC3 correspond to the three different scenarios, where the two groups exhibit spatio-temporal differences (SC1), temporal differences (SC2), and spatial differences (SC3), respectively.

scenario (SC3), the two groups only differ in terms of the spatial locations of the voxels, which have been assigned the second temporal profile.

In addition, we also studied the effect of the signal-to-noise ratio (SNR) on the performance of our inferential methods. In particular, we varied our choice of σ , when generating the different time series displayed in panels (d-f) of figure 1, in order to produce different SNRs. In these simulations, the ‘signal’ of interest was defined as the amplitude of the original sinusoids, which oscillated between -1 and 1, thus giving an amplitude of $\lambda = 2$. Since the noise affecting this signal was specified to be Gaussian, the SNR was defined, in our case, as $\text{SNR} = \lambda/2\sigma$. Thus, by setting σ to either 1/2, 1 or 2, we produced three different SNRs of 2, 1 and 1/2, respectively.

These synthetic data sets were analyzed using our proposed inferential framework. For each simulated subject-specific volume, a SOM was computed, and the restricted Frechet mean was identified for each group. In all scenarios, SOMs were produced by using the batch SOM algorithm. The output grid was of size 3×3 with $K = 9$; the number of iterations was set to 100 steps; and we used a decreasing neighborhood kernel of size $k_1 = 3$, as commonly done in this field (Kohonen et al., 2000). For computational convenience, statistical inference was drawn in each scenario after 100 permutations. Each simulated scenario was reproduced 100 times, and constructed for the three different levels of SNR, thereby totalling 900 distinct simulations.

Simulation Results

The summary results of the analysis of these synthetic data sets are reported in table 1 and figure 2. Overall, the different metrics of interest were found to successfully capture the aspects of the simulated SOMs that they were expected to identify. That is, in the first column of table 1, one can see that T-SMD, which solely takes into account the differences in voxel-specific temporal profiles, attains its most significant values under the temporal scenario, SC2. Similarly, in the second column of table 1, the spatial version of the SMD metric, denoted S-SMD exhibits its best performance under the first and third scenarios, denoted SC1 and SC3, respectively. Indeed, these two scenarios are the only ones, where the two groups can be discriminated in terms of the spatial locations of the different types of time series. Finally, in the third column of table 1, the

<i>Scenarios and Factors</i>		T-SMD	S-SMD	ST-SMD
SC1 (Spatio-temporal)	SNR = 2	0 ± 0	0.012 ± 0.033	0 ± 0
	SNR = 1	0 ± 0.001	0.518 ± 0.171	0.003 ± 0.012
	SNR = 0.5	0.030 ± 0.066	0.800 ± 0.235	0.049 ± 0.123
SC2 (Temporal)	SNR = 2	0 ± 0	0.499 ± 0.303	0 ± 0.006
	SNR = 1	0 ± 0	0.499 ± 0.295	0.001 ± 0.005
	SNR = 0.5	0.017 ± 0.070	0.484 ± 0.296	0.022 ± 0.030
SC3 (Spatial)	SNR = 2	0.472 ± 0.294	0.014 ± 0.057	0.029 ± 0.055
	SNR = 1	0.464 ± 0.286	0.525 ± 0.167	0.109 ± 0.122
	SNR = 0.5	0.525 ± 0.279	0.783 ± 0.271	0.101 ± 0.141

Table 1. Significance levels based on synthetic data with 100 simulations in every cell, with the mean p -value and the standard deviation for that distribution of p -values. These results are reported for the three scenarios described in figure 1, which are denoted by SC1, SC2 and SC3, for three different levels of SNR, and for the three different distance functions under scrutiny, denoted by T-SMD, S-SMD and ST-SMD, which stand for the temporal SMD, spatial SMD, and spatio-temporal SMD, respectively.

spatio-temporal metric, ST-SMD, appears to be optimal under the first scenario, where group differences can be characterized through the spatio-temporal properties of the simulated images.

In addition, we have also evaluated the effect of sample size on the capacity of the metrics to detect group differences. These results are not reported in this paper, but we have observed, as expected, that the statistical power of all the studied metrics improved as the number of subjects in each group increases. In particular, it was noted that for the ST-SMD, we solely needed $n \geq 15$ in each group to identify significant differences under SNR=1, and group sizes of $n \geq 20$ under SNR=0.5; for all scenarios. Exemplary null distributions for t_F -statistics in the three different scenarios and the three different levels of SNR with $n = 15$ are reported in figure 2.

In sum, one can note that these three distance functions exhibit different levels of robustness. In particular, S-SMD appeared to be especially sensitive to noise. Although S-SMD succeeded to capture the spatial differences simulated in SC3, it only outperformed ST-SMD for high SNRs. Also, in the spatio-temporal scenario (SC1), T-SMD behaved as well or better than ST-SMD. This suggests that the T-SMD function is more ‘powerful’ than the ST-SMD, even when the group differences are characterized by both spatial and temporal properties. However, the use of ST-SMD remains justified, because it also succeeds to capture spatial differences, whereas T-SMD fails to do so. In general, we therefore recommend the joint use of T-SMD and ST-SMD: If only T-SMD indicates the presence of group differences, then one can conclude that such differences are mainly of a temporal nature; whereas when the use of ST-SMD indicates greater group differences, this suggests that such differences also have a spatial character. Overall, these simulations highlight the importance of using several types of distance functions, as there may not exist a single type of metric, which would capture all of the aspects of the data of interest.

Experimental Data

We also evaluated our methods with the re-analysis of a classical data set, originally published by Mourao-Miranda et al. (2006). Since this first publication, this data set has been re-analyzed several times with different machine learning algorithms, as conducted by Mourao-Miranda et al. (2007) using spatio-temporal support vector machine (SVM), and Hardoon et al. (2007) with unsupervised methods.

Subjects and Experimental Design

This data set consists of fMRI data from 16 right-handed males with a mean age of 23 years. All participants had normal eyesight and no history of neurological or psychiatric disorders, and gave written informed consent to participate in the study, in accordance with the local ethics




Figure 2. Histograms of the null distributions of t_F -values obtained through permutation. These null distributions are given for a single synthetic data set under the three different simulation scenarios, denoted SC1, SC2 and SC3, respectively, and for three different metrics on the space of SOMs, denoted T-SMD, S-SMD and ST-SMD, which stand for sums of minimum distances, spatial T-SMD and spatio-temporal T-SMD, respectively. The red dashed line indicates the value of the actual t_F -statistic for the simulation of interest. These histograms were constructed using data based on an SNR of 1, and for 15 subjects in each group.

<i>Metrics</i>	<i>Tests</i>	<i>p-values</i>
T-SMD	Pleasant vs. Neutral	0.0
	Unpleasant vs. Neutral	0.0
	Pleasant vs. Unpleasant	0.258
S-SMD	Pleasant vs. Neutral	0.412
	Unpleasant vs. Neutral	0.518
	Pleasant vs. Unpleasant	0.423
ST-SMD	Pleasant vs. Neutral	0.021
	Unpleasant vs. Neutral	0.013
	Pleasant vs. Unpleasant	0.128

Table 2. Significance results of all pairwise comparisons for the three conditions of interest, where subjects were exposed to pleasant, unpleasant and neutral stimuli. These results are reported independently for three different metrics, denoted T-SMD, and S-SMD and ST-SMD, which stand for temporal SMD, spatial SMD and spatio-temporal SMD, respectively. Observe that since three tests were conducted for each pair of conditions and we therefore necessitate the application of a Bonferroni correction for testing for these three pairwise comparisons. Hence, only p -values below 0.016 should be regarded as statistically significant.

committee of the University of North Carolina (see Mourao-Miranda et al., 2006). Data were acquired using an experimental block design, composed of three different conditions: (i) exposure to unpleasant visual stimuli (i.e. photos of dermatological diseases), (ii) exposure to neutral visual stimuli (i.e. photos of neutral day-to-day scenes including human actors) and (iii) exposure to male-relevant pleasant visual stimuli (i.e. scantily dressed women or women in swimsuits). The entire experimental design consisted of six blocks, where each block contained seven images, which were each presented to the subjects for three seconds. Each block was followed by a resting block period where subjects were solely exposed to a fixation cross. All blocks were of 21s in length.

Data Acquisition and Pre-Processing

Blood-oxygenation-level-dependent (BOLD) signal was measured using a 3-Tesla Allegra head-only MRI System at the Magnetic Resonance Imaging Research Center in the University of North Carolina. The scanning parameters were specified as follows. Voxel size was $3 \times 3 \times 3mm^3$, TR was 3s, TE was 30ms, FA was 80, FOV was $192 \times 192mm$ and each MRI volume had dimensions $64 \times 64 \times 49$. In each experiment, a total of 254 functional volumes were acquired for each participant.

Data were pre-processed using the FSL Software suite (Smith et al., 2004a); through the use of the Nipype Python Library (Gorgolewski et al., 2011). All fMRI volumes were first motion-corrected and the skulls were removed, after tissue segmentation. The voxel time series were detrended and filtered in time and space: that is, low-frequency (drift) fluctuations were reduced using a band-pass temporal filter comprised between 0.008Hz and 0.1Hz. The use of a such a band-pass filter is typical of resting-state connectivity analysis (Cordes et al., 2001). In addition, spatial smoothing was performed using an 8mm full-width at half-maximum Gaussian kernel.

The first two volumes of each block were discarded, to allow for the between-block lag in haemodynamic response. The remaining volumes were concatenated to form three distinct time series representing the three different conditions. Time series concatenation in the context of functional connectivity has been introduced by Fair et al. (2007) and has been implemented by several authors for the study of functional MRI networks (Ginestet and Simmons, 2011, Ginestet et al., 2011).

Results of Real-data Analysis

The significance results of all the pairwise comparisons are reported in table 2, after applying the Bonferroni correction for multiple testing. Our re-analysis of this data set has highlighted a

<i>Ranked SOM Units</i>	<i>Sample Jaccard Index</i>		
	<i>Neutral</i>	<i>Pleasant</i>	<i>Unpleasant</i>
Units up to 1	0.68	0.32	0.55
Units up to 2	0.11	0.28	0.16
Units up to 3	0.08	0.04	0.03
Units up to 4	0.01	0.07	0.07
Units up to 5	0.03	0.08	0.03
Units up to 6	0.01	0.05	0.04
Units up to 7	0.05	0.08	0.05
Units up to 8	0.01	0.02	0.05
Units up to 9	0.02	0.06	0.02

Table 3. Individual percentage overlaps between the group-level SOM components and thresholded GLM z -score maps, where the SOM components have been ranked with respect to the sample Jaccard index. In bold, we have highlighted the three ‘best’ condition-specific SOM output units based on T-SMD, which are visually described in figures 3 to 5. Each column sums to 1.0, since the concatenated SOM output units cover all the voxels in the fMRI volumes of interest.

substantial degree of difference between the neutral and pleasant conditions, on one hand, and between the neutral and unpleasant conditions, on the other hand. These pairwise differences were found to be highly significant under the T-SMD distance function. The mean SOM in the unpleasant condition was also found to be significantly different from the mean SOM in the neutral condition with respect to the ST-SMD metric ($p < 0.016$), albeit to a lesser extent than under the T-SMD function. The difference between the mean SOMs in the pleasant condition and the neutral one was found to be just above significance level under the ST-SMD metric ($p = 0.063$). By contrast, none of these differences reached significance under the S-SMD, indicating that differences in spatial allocation of the different output units of the SOMs in these experimental conditions were not sufficient to distinguish between the group-level SOMs. Importantly, the mean SOMs under the pleasant and unpleasant conditions were not found to be significantly different under all three metrics.

As noted in the analysis of the synthetic data, the fact that the SOMs are computed on the basis of the similarities between the voxel-specific time series is likely to be responsible for the important differences that we are reporting between the metrics capturing the temporal aspects of the process and S-SMD, which does not emphasize differences in temporal profiles. Intriguingly, it should also be noted that the differences between the SOMs in the pleasant and unpleasant conditions under ST-SMD is lower than under the T-SMD, thereby perhaps suggesting that the differences between these two conditions is characterized by an ‘interaction’ between the temporal and spatial properties of these functional patterns.

Visualization of Group-level SOMs

In each of the three conditions, we have represented the subject-specific restricted Frechet mean in order to produce robust visual summaries of the different output units in each condition. This was conducted by identifying the output units that explained the largest amount of ‘variance’ in the data, in terms of *sample Jaccard index*. This measure quantifies how representative is a mean output unit in terms of the overlap of this unit with the output units of all the subjects in that group.

For each experimental condition, we identified the output unit in the subject-specific SOMs that explained the largest amount of Jaccard index. For each experimental condition, we have plotted in figures 3, 4 and 5, the three output units that are associated with the least amount of Jaccard index over all subjects. That is, for the j^{th} condition, the sample Jaccard index of the




Figure 3. Representation of three output units of the restricted Frechet mean SOM for the *unpleasant* condition in red, with thresholded ($p \leq 0.05$) GLM z -score maps in blue. The output units have been projected in MNI-normalized space. These three output units are the ones that explain the largest amount of sample Jaccard index in that SOM. They are ordered by Jaccard index from panels (a) to (c), with the unit exhibiting the smallest Jaccard index in (a).

k^{th} output unit in the restricted Frechet mean for that condition was defined as follows,

$$J(\mathbf{m}_{kj}) = \frac{1}{n_j} \sum_{i=1}^{n_j} \mathbf{J}^k(\mathbf{M}_j, \mathbf{X}_{ij}),$$

Here, $\mathbf{J}^k(\mathbf{M}, \mathbf{X}) = \sum_{v=1}^V \text{Jacc}\{S(\mathbf{x}_v), S(\mathbf{w}_k)\}$ denotes the *global* Jaccard distance between a mean SOM and a subject-specific SOM. Moreover, the classical vector-specific Jaccard distance is

$$\text{Jacc}\{\mathbf{x}, \mathbf{y}\} = \frac{C_{10} + C_{01}}{C_{11} + C_{01} + C_{10}},$$

where C_{01} is the number of elements satisfying $x_h = 1$ and $y_h = 0$, for h running from 1 to the length of these vectors; and C_{10} and C_{11} are defined similarly. The three output units minimizing the sample Jaccard index in each of the three experimental conditions have been plotted in figures 3, 4 and 5, for the neutral, pleasant, and unpleasant conditions, respectively. Allocation of a voxel to the particular output unit of a SOM is ‘hard’, in the sense, that either a voxel is included into an output unit or it is included in a different one. Thus, in figures 3, 4 and 5, we have provided the spatial location of the voxels that have been assigned to particular output units in each condition.

From these three figures, one can observe that the neutral condition is characterized by a distinct network of brain regions identified as the third output unit in the neutral condition, as can be seen from panel (c) of figure 3. The set of regions associated with this particular output unit may be interpreted as a visual network, since it contains a considerable number of regions located in the occipital lobe. This particular output unit was not found to be present in the three units with the least Jaccard index in either the pleasant or unpleasant condition, as can be noted from figures 4 and 5, respectively.

Comparison with Standard GLM Maps

The group-level SOM output units selected using the sample Jaccard index were compared with standard general linear model (GLM) z -score maps. Separate GLM analyses were conducted for each experimental condition, using the FEAT function provided in the FSL software suite (Smith et al., 2004b). The z -score maps were thresholded at $p = 0.05$, for comparison purposes. These binarized z -score maps were compared with the maps of the three ‘best’ group-level output units, selected on the basis of their sample Jaccard indices, as described in the previous section. These




Figure 4. Representation of three output units of the restricted Frechet mean SOM for the *unpleasant* condition in red, with thresholded ($p \leq 0.05$) GLM z -score maps in blue. The output units have been projected in MNI-normalized space. These three output units are the ones that explain the largest amount of sample Jaccard index in that SOM. They are ordered by Jaccard index from panels (a) to (c), with the unit exhibiting the smallest Jaccard index in (a).

thresholded GLM maps have been overlaid in blue over the selected output units in figures 3, 4 and 5.

In each experimental condition, we computed the percentage overlap between the maps obtained using these two different techniques. That is, in each condition, we evaluated how many voxels were present in both the thresholded z -score maps and each output unit, normalized by the number of voxels included in the z -score maps. These numerical comparisons are reported in table 3, where we have described the individual percentage overlap of the nine different output units, ranked with respect to their sample Jaccard indices. The combined percentage overlap of the three output units exhibiting the lowest sample Jaccard indices with the thresholded GLM z -score maps was 87%, 64%, and 74% for the neutral, pleasant and unpleasant conditions, respectively. Although our proposed SOM-based methods summarize such fMRI volumes in a non-linear manner, these numerical comparisons show that the resulting output units exhibit a considerable degree of agreement with standard GLM analysis.

Discussion

Advantages of Proposed Methods

The main contribution of this paper is the construction of an inferential framework for the comparison of group-level SOMs. Although some previous researchers have considered various ways of comparing SOMs (Kaski and Lagus, 1996, Deng, 2007, Kirt et al., 2007), to the best of our knowledge, no authors have yet treated the problem of evaluating the statistical significance of such group differences. In addition, observe that our Frechean approach to statistical inference can be conducted for any choice of metrics. Indeed, the idea of defining a group distance statistic, such as the Frechean t -test in equation (9) and then evaluating its significance using permutation can be implemented for any choice of distance functions. Therefore, this allows the specification of a rich array of different distance functions capturing different aspects of the SOMs under scrutiny. As illustrated in the main body of the paper, we have shown that classical distance functions such as the SMD can be modified in order to emphasize spatial, temporal or spatio-temporal differences between the groups of interest. Here, the choice of hypothesis is therefore superseded by the choice of metrics over the space of SOMs. In particular, this inferential framework allows to test hitherto untestable group-level hypotheses.

Another substantial advantage of combining a Frechean approach with the computation of




Figure 5. Representation of three output units of the restricted Frechet mean SOM for the *unpleasant* condition in red, with thresholded ($p \leq 0.05$) GLM z -score maps in blue. The output units have been projected in MNI-normalized space. These three output units are the ones that explain the largest amount of sample Jaccard index in that SOM. They are ordered by Jaccard index from panels (a) to (c), with the unit exhibiting the smallest Jaccard index in (a).

subject-specific SOMs is that this bypasses the problem of multiple testing correction. In standard mass-univariate analyses of MRI volumes, one needs to control for the inflation of the number of false positives introduced by performing a large number of non-independent statistical tests. By contrast, we are here conducting a single test, which identifies whether the volumes of interest are different at a multivariate level, through the comparison of two non-parametric unsupervised representations of the original data.

Finally, one should also note that the use of the *restricted* Frechet mean in our proposed framework is advantageous for several reasons. On the one hand, the restricted Frechet mean greatly reduces the computational cost of our overall analytical procedure. This is especially true, because inference was drawn using permutations of the group labels, and it is not clear whether such a large number of permutations would have been possible, lest for the use of the restricted Frechet mean. On the other hand, the restricted Frechet mean has also the advantage of quasi-automatically transforming any distance function into a proper metric that satisfies the four metric axioms. This results in a non-negligible simplification of the probabilistic theory needed to justify our inferential approach. Indeed, most of the asymptotic results, which have been previously established relies on the postulate that the distance function of interest is a proper metric (Ziezold, 1977, Sverdrup-Thygeson, 1981).

Limitations of the Frechean Framework

One can identify three substantial limitations to our proposed Frechean inferential framework for SOMs, which are (i) a lack of contrast maps, (ii) a reliance on permutation for statistical inference, and (iii) the use of the *restricted* Frechet mean in the place of the unrestricted mean element in the space of all SOMs. We address these three limitations in turn.

Firstly, one of the important limitations of our current method is that it does not directly permit the production of a ‘group-difference SOM’ representing the difference between two group-specific Frechet means. In particular, this implies that we cannot represent such differences by plotting a differential pattern of activation or connectivity, as is commonly done using standard mass-univariate approaches. See the statistical parametric network (SPN) approach, advocated by Ginestet and Simmons (2011) for example, when conducting functional network analyses. From a neuroscientific perspective, this is a considerable limitation, as it diminishes the interpretability of the results. We will consider different ways of tackling this issue and producing group-difference SOMs in future work.

Secondly, our inferential framework has relied on permutation for evaluating the statistical significance of the test statistics under scrutiny. This was made computationally feasible, because this choice of inferential method was used in conjunction with the restricted Frechet mean. That is, for each permutation of the group labels, the computation of the group-specific Frechet means was straightforward, because the identification of the restricted Frechet mean can be conducted by using the margins of the dissimilarity matrix of the sample points –that is, the dissimilarity matrix of all the subject-specific SOMs. Hence, the cost of calculating the group means at each permutation was small, and the full inference could be drawn within a couple of hours on a standard desktop computer. Such a level of computational efficiency may not have been achieved if we had attempted to derive the unrestricted Frechet mean, as described in equation (6), which would have necessitated to perform a minimization over the space of all possible SOMs.

However, although the use of the restricted Frechet mean was advantageous from a computational perspective, this particular methodological choice has also its limitations. Indeed, the use of the restricted Frechet mean in the place of the unrestricted mean results in a loss of the classical benefits usually associated with computing an average of real numbers. In decision-theoretic parlance and when considering real-valued random variables, the arithmetic mean is the quantity that minimizes the squared error loss (SEL) (see Berger, 1980, for an introduction to decision theory). The restricted version of the arithmetic mean for real-valued random variables would also minimize the restricted SEL. However, the restricted arithmetic mean would necessarily achieve a sample variance greater or equal to the one of the unrestricted Frechet mean. Note, however, that the problems associated with the utilization of the restricted Frechet mean are also mitigated by the fact that computing this quantity quasi-automatically makes the space of interest a metric space, regardless of the particular choice of distance function.

We have here used the sample Jaccard index for selecting the most ‘relevant’ output units in the group-level SOMs. It certainly does not follow from such a selection criterion that these output units are of greater physiological relevance. This criterion is entirely statistical, and the resulting interpretation of these output units should remain statistical. In practice, it is advisable to visualize the entire set of output units obtained after this type of SOM analysis, in order to identify relevant physiological differences based on prior neuroanatomical knowledge.

Possible Extensions of these Methods

Our proposed Frechean inferential framework could be extended in a range of different directions. One of the most natural extensions of this method would be to devise an F -test, which would generalize the aforementioned two-sample t_F -statistic. A Frechet F -statistic may take the following form. Let a data set of the form $\mathbf{M}_{ij} \in (\mathcal{M}, d)$, where $i = 1, \dots, n_j$ labels the objects in the j^{th} group with $j = 1, \dots, J$. By analogy with the classical real-valued setting, the F -statistic can be defined as the ratio of the between-group to within-group variances, $F_F = SS_1 / SS_0$, where these quantities are here defined with respect to the Frechet moments, such that $SS_1 = (J - 1)^{-1} \sum_{j=1}^J n_j d(\overline{\mathbf{M}}_j, \overline{\mathbf{M}})^2$, and $SS_0 = (N - J)^{-1} \sum_{j=1}^J \sum_{i=1}^{n_j} d(\mathbf{M}_{ij}, \overline{\mathbf{M}}_j)^2$, using standard notation for the Frechet sample group means, $\overline{\mathbf{M}}_j$, and grand mean, $\overline{\mathbf{M}}$. One can then test for the null hypothesis that $H_0 : \sigma_1^2 = \sigma_0^2$, where σ_1^2 and σ_0^2 are the theoretical equivalents of SS_1 and SS_0 , respectively. Statistical inference can, again, be conducted using permutation of the group labels.

In addition, the analytical strategy that we have here described could also be improved through the use of different types of SOM algorithm. In the present paper, we have made use of the batch SOM algorithm. However, several other alternatives to the traditional sequential SOM algorithm have been proposed in the literature. In particular, Vesanto and Alhoniemi (2000) have showed that the SOMs obtained when using the batch SOM algorithm with an initialization of the maps based on the eigenvectors of the input data can produce more robust results. Since every SOM is computed independently for each subject, such an improvement of the existing batch SOM algorithm could easily be incorporated in our proposed inferential framework.

One of the outstanding questions that is implicitly raised in this paper is the possibility of sep-

arately weighting the individual contributions of the spatial and temporal properties contributing to the overall SOM difference. Such a question is likely to be arduous to answer, however, since the temporal and spatial properties of the fMRI volumes of interest necessarily live in distinct abstract spaces. On one hand, the temporal differences in T-SMD were quantified using a Euclidean distance in a T -dimensional vector space; whereas, on the other hand, the spatial differences in S-SMD were quantified using the Hamming distance on binary vectors of varying sizes. It is unclear whether the magnitude of the distances in these different metric spaces could be normalized in order to ensure a modicum of comparability.

Conclusions

In this paper, we have described a formal framework for drawing group-level inference between unsupervised multivariate summaries of fMRI data. Our proposed approach proceeds by computing subject-specific SOMs, and computing the sample Frechet mean in each group of subjects. Despite the unwieldy nature of the space of all possible SOMs, this can be done efficiently by identifying the restricted Frechet mean. Statistical inference on the difference between the group restricted Frechet means can be conducted using permutation on the group labels. This framework can be implemented for any choice of metrics quantifying the difference between pairs of SOMs. As such, the specification of a particular distance function is equivalent to the choice of a particular hypothesis test. Different researchers may therefore be interested in evaluating different metrics, which capture different aspects of the SOMs.

We have hence described and evaluated several types of distance functions for SOMs based on fMRI data. In particular, we have considered variants of the classic SMD function, which has previously been used to compare pairs of SOMs. Our proposed variants distinguish between the temporal, spatial and spatio-temporal properties of the data under scrutiny. Our inferential framework and these metrics were tested on both synthetic and real data. Our analysis of the simulated data showed that the distance functions of interest were indeed capturing the aspects of the data that they were purported to measure. In addition, the findings of the re-analysis of an fMRI experiment has demonstrated the capacity of these methods to extract new information from existing data sets. In this paradigm, the differences of the restricted mean SOMs in the pleasant and unpleasant conditions were found to be smaller than the differences between the mean SOMs in any of these two conditions with respect to the one in the neutral condition.

Taken together, the analyses of these synthetic and real data sets have underlined the robustness and potential usefulness of these methods. It is hoped that this type of global inferential perspective on neuroimaging data will inspire other neuroscientists to follow this research avenue. One could imagine a range of other subject-specific abstract-valued random variables that could be suitably analyzed using this type of Frechet inferential framework. In fact, the very use of mass-univariate approaches in the context of neuroimaging could be superseded by a more global perspective, where a single statistical test is conducted, thereby bypassing the need for exacting multiple testing penalties.

Appendices

A. From Distance Functions to Metrics

Let d be a distance function on a *finite* space of SOMs, $\mathbf{\Lambda} = \{\mathbf{M}_1, \dots, \mathbf{M}_n\}$, which satisfies the positivity, coincidence and symmetry axioms. In order to transform the distance function d into a proper metric \tilde{d} satisfying the triangle inequality, we need to construct a saturated graph $G = (V, E)$ representing the topology of $\mathbf{\Lambda}$. The vertex set of G is defined as $V(G) = \mathbf{\Lambda}$. Its edge set is composed of all the possible links between the elements of $\mathbf{\Lambda}$. That is, G is a saturated graph, in the sense that it contains the maximal number of edges. Each of these edges is denoted by $\mathbf{M}_i \mathbf{M}_j \in E(G)$, for any $0 \leq i \neq j \leq n$.

<i>Scenarios</i>	<i>SOM Dimensions</i>	T-SMD	S-SMD	ST-SMD
SC1 (Spatio-temporal)	10×10	0 ± 0	0.626 ± 0.236	0.002 ± 0.064
	5×5	0 ± 0	0.498 ± 0.256	0.001 ± 0.031
	4×6	0 ± 0	0.516 ± 0.285	0.001 ± 0.012
	6×8	0 ± 0	0.464 ± 0.279	0.001 ± 0.078
SC2 (Temporal)	10×10	0 ± 0	0.612 ± 0.350	0 ± 0
	5×5	0 ± 0	0.557 ± 0.298	0.011 ± 0.014
	4×6	0 ± 0	0.474 ± 0.288	0.002 ± 0.012
	6×8	0 ± 0.006	0.487 ± 0.269	0.003 ± 0.097
SC3 (Spatial)	10×10	0.505 ± 0.296	0.523 ± 0.282	0.180 ± 0.171
	5×5	0.519 ± 0.206	0.504 ± 0.279	0.149 ± 0.144
	4×6	0.482 ± 0.272	0.559 ± 0.268	0.108 ± 0.162
	6×8	0.482 ± 0.300	0.451 ± 0.281	0.149 ± 0.103

Table 4. Simulation results with varying SOM dimensions summarized as mean significance levels and standard deviations of these distributions, based on synthetic data with 100 simulations in every cell and $\text{SNR} = 1$. These results are reported for the three scenarios described in figure 1, which are denoted by SC1, SC2 and SC3, for three different levels of SNR, and for the three different distance functions under scrutiny, denoted by T-SMD, S-SMD and ST-SMD, which stand for the temporal SMD, spatial SMD, and spatio-temporal SMD, respectively. These results are consistent with the ones of table 1.

A *path* in G is a non-empty subgraph $P \subseteq G$ of the form $V(P) = \{\mathbf{M}_0, \dots, \mathbf{M}_k\}$ and $E(P) = \{\mathbf{M}_0\mathbf{M}_1, \dots, \mathbf{M}_{k-1}\mathbf{M}_k\}$, where the \mathbf{M}_i 's are all *distinct*. Following an idea proposed by Mannila and Eiter (1997), it is now possible to construct a new distance function, denoted \tilde{d} , defined as the set of *shortest paths* in Λ , such that for any $\mathbf{M}, \mathbf{M}' \in \Lambda$, we have

$$\tilde{d}(\mathbf{M}, \mathbf{M}') = \min_{P \in \mathcal{P}(\mathbf{M}, \mathbf{M}')} \sum_{\mathbf{M}_i \mathbf{M}_j \in E(P)} d(\mathbf{M}_i, \mathbf{M}_j),$$

where $\mathcal{P}(\mathbf{M}, \mathbf{M}')$ is the set of all paths in G between \mathbf{M} and \mathbf{M}' . By construction, it immediately follows that \tilde{d} satisfies the triangle inequality. Therefore, (Λ, \tilde{d}) forms a proper metric space.

B. Choice of SOM Dimensions

A supplemental set of simulations was conducted in order to investigate the effect of the choice of SOM dimensions on group-level statistical inference. We assessed the effect of rectangular SOMs, as well as the effect of increasing the dimensions of these maps. The synthetic data used for these simulations followed the design described in the section entitled Synthetic Data Simulations, based on the three different scenarios, and using the three types of SMD functions described in this paper, and setting $\text{SNR} = 1$. These results are reported in table 4.

The results of these simulations were consistent with the ones described in our first analysis of these synthetic data. In particular, for any choice of SOM dimensions, we obtained strong corroborations of the previous findings. Under both SC1 and SC2, the T-SMD tended to outperform its counterparts for any choice of SOM dimensions. As before, S-SMD performed poorly throughout these simulations, irrespective of the choice of SOM dimensions. Finally, the ST-SMD function exhibited good performance on all scenarios, and outperformed the T-SMD under SC3, although ST-SMD did not reach significance level for this particular scenario.

References

- Balding, D., Ferrari, P., Fraiman, R., and Sued, M. (2009). Limit theorems for sequences of random trees. *TEST*, **18**(2), 302–315.
- Berger, J. (1980). A Robust Generalized Bayes Estimator and Confidence Region for a Multivariate Normal Mean. *The Annals of Statistics*, **8**(4), 716–761.

- Bigot, J. and Charlier, B. (2011). On the consistency of Frechet means in deformable models for curve and image analysis. *Electronic Journal of Statistics*, **5**, 1054–1089.
- Cordes, D., Haughton, V., Arfanakis, K., Carew, J., Turski, P., Moritz, C., Quigley, M., and Meyerand, M. (2001). Frequencies contributing to functional connectivity in the cerebral cortex in resting-state data. *American Journal of Neuroradiology*, **22**(7), 1326–1333.
- Deng, D. (2007). Content-based image collection summarization and comparison using self-organizing maps. *Pattern Recognition*, **40**(2), 718–727.
- Dryden, I. and Mardia, K. (1998). *Statistical Analysis of Shape*. Wiley, London.
- Fair, D., Schlaggar, B., Cohen, A., Miezin, F., Dosenbach, N., Wenger, K., Fox, M., Snyder, A., Raichle, M., and Petersen, S. (2007). A method for using blocked and event-related fMRI data to study “resting state” functional connectivity. *NeuroImage*, **35**(1), 396–405.
- Fischer, H. and Hennig, J. (1999). Neural network-based analysis of MR time series. *Magnetic resonance in medicine*, **41**(1), 124–131.
- Fréchet, M. (1948). Les éléments aléatoires de nature quelconque dans un espace distancié. *Annales de L’Institut Henri Poincaré*, **10**(4), 215–310.
- Ginestet, C., Nichols, T., Bullmore, E., and Simmons, A. (2011). Brain network analysis: Separating differences in cost from differences in topology. *PLoS ONE*, **6**(7), e21570.
- Ginestet, C. and Simmons, A. (2011). Statistical parametric network analysis of functional connectivity dynamics during a working memory task. *NeuroImage*, **5**(2), 688–704.
- González, F. and Dasgupta, D. (2003). Anomaly detection using real-valued negative selection. *Genetic Programming and Evolvable Machines*, **4**(4), 383–403.
- Gorgolewski, K., Burns, C.D., Madison, C., Clark, D., Halchenko, Y.O., Waskom, M.L., and Ghosh, S.S. (2011). Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python. *Front Neuroinform*, **5**.
- Hamming, R. (1950). Error detecting and error correcting codes. *Bell System Technical Journal*, **29**, 147–160.
- Hardoon, D., Mourão Miranda, J., Brammer, M., and Shawe-Taylor, J. (2007). Unsupervised analysis of fmri data using kernel canonical correlation. *NeuroImage*, **37**(4), 1250–1259.
- Karcher, H. (1977). Riemannian center of mass and mollifier smoothing. *Communication in Pure and Applied Mathematics*, **30**, 509–541.
- Kaski, S. and Lagus, K. (1996). Comparing self-organizing maps. *Artificial Neural Networks—ICANN 96*, 809–814.
- Kirt, T., Vainik, E., and Vöhandu, L. (2007). A method for comparing self-organizing maps: case studies of banking and linguistic data. In *Eleventh East-European Conference on Advances in Databases and Information Systems ADBIS, Varna, Bulgaria, Technical University of Varna*, 107–115.
- Kohonen, T. (2001). *Self-Organizing Maps*. Springer Series in Information Sciences. Springer.
- Kohonen, T., Kaski, S., Lagus, K., Salojärvi, J., Honkela, J., Paatero, V., and Saarela, A. (2000). Self organization of a massive document collection. *IEEE Transactions on Neural Networks*, **11**(3), 574–585.
- Kriegel, H. (2004). Classification of websites as sets of feature vectors. *Proc. IASTED DBA*, 127–132.
- Liao, W., Chen, H., Yang, Q., and Lei, X. (2008). Analysis of fMRI data using improved self-organizing mapping and spatio-temporal metric hierarchical clustering. *IEEE Transactions on Medical Imaging*, **27**(10), 1472–1483.
- Mannila, H. and Eiter, T. (1997). Distance Measures for Point Sets and Their Computation. *Acta Informatica*, **34**, 109–133.

- Mourao-Miranda, J., Friston, K., and Brammer, M. (2007). Dynamic discrimination analysis: a spatio-temporal SVM. *NeuroImage*, **36**(1), 88–99.
- Mourao-Miranda, J., Reynaud, E., McGlone, F., Calvert, G., and Brammer, M. (2006). The impact of temporal compression and space selection on svm analysis of single-subject and multi-subject fmri data. *NeuroImage*, **33**(4), 1055–1065.
- Ngan, S.C. and Hu, X. (1999). Analysis of functional magnetic resonance imaging data using self-organizing mapping with spatial connectivity. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, **41**(5), 939–946.
- Ngan, S.C., Yacoub, E.S., Auffermann, W.F., and Hu, X. (2002). Node merging in Kohonen’s self-organizing mapping of fMRI data. *Artificial intelligence in medicine*, **25**(1), 19–33.
- Parthasarathy, K. (1967). *Probability Measures on Metric Spaces*. American Mathematical Society, London.
- Peltier, S.J., Polk, T.A., and Noll, D.C. (2003). Detecting low-frequency functional connectivity in fMRI using a self-organizing map (SOM) algorithm. *Human Brain Mapping*, **20**(4), 220–226.
- Searcóid, M. (2007). *Metric spaces*. Springer, London.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., and Matthews, P.M. (2004a). Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, **23** Suppl 1, S208–19.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., and Matthews, P.M. (2004b). Advances in functional and structural mr image analysis and implementation as fsl. *NeuroImage*, **23**(Supplement 1), S208–S219.
- Sverdrup-Thygeson, H. (1981). Strong law of large numbers for measures of central tendency and dispersion of random variables in compact metric spaces. *The Annals of Statistics*, **9**(1), 141–145.
- Takala, V., Ahonen, T., and Pietikainen, M. (2005). Block-based methods for image retrieval using local binary patterns. In *Proc. 14th Scandinavian Conference on Image Analysis (SCIA)*, 882–891.
- Tan, X., Chen, S., Zhou, Z., and Zhang, F. (2005). Recognizing partially occluded, expression variant faces from single training image per person with som and soft k-nn ensemble. *Neural Networks, IEEE Transactions on*, **16**(4), 875–886.
- Tarca, A.L., Carey, V.J., Chen, X.w., Romero, R., and Drghici, S. (2007). Machine learning and its applications to biology. *PLoS Comput Biol*, **3**(6), e116.
- Thorstensen, N., Segonne, F., and Keriven, R. (2009). *Scale Space and Variational Methods in Computer Vision*, vol. 5567, chap. Pre-image as Karcher Mean Using Diffusion Maps: Application to Shape and Image Denoising, 721–732. Springer.
- Tungaraza, R., Guan, J., Rolfe, S., Atmosukarto, I., Poliakov, A., Kleinhans, N., Aylward, E., Ojemann, J., Brinkley, J., and Shapiro, L. (2009). A similarity retrieval method for functional magnetic resonance imaging (fmri) statistical maps. In *Proceedings of SPIE*, vol. 7259, 72590E.
- Vesanto, J. and Alhoniemi, E. (2000). Clustering of the self-organizing map. *Neural Networks, IEEE Transactions on*, **11**(3), 586–600.
- Wismüller, A., Meyer-Bäse, A., Lange, O., Auer, D., Reiser, M., and Sumners, D. (2004). Model-free functional mri analysis based on unsupervised clustering. *Journal of Biomedical Informatics*, **37**(1), 10–18.
- Ziezold, H. (1977). On expected figures and a strong law of large numbers for random elements in quasi-metric spaces. *Transactions of the Seventh Prague Conference on Information Theory, Statistical Decision Functions, Random Processes and of the 1974 European Meeting of Statisticians*.